



Defence Research and
Development Canada Recherche et développement
pour la défense Canada



Accelerated training for command dynamic decision making:

A Pilot Study Using Microworlds

Jerzy Jarmasz

Defence R&D Canada
Technical Report
DRDC Toronto TR 2006-239
December 2006

Canada

Accelerated Training for Command Dynamic Decision Making

A Pilot Study Using Microworlds

Jerzy Jarmasz

Defence R&D Canada – Toronto

Technical Report

DRDC Toronto TR 2006-239

December 2006

Principal Author

Original approved by

Jerzy Jarmasz

Ph.D.

Approved by

Original approved by

Joseph V. Baranski

Head, Collaborative Behaviour and Integrated Learning Section

Approved for release by

Original approved by

K.M. Sutton

Chair, Document Review and Library Committee

The work reported in this document was performed under Thrust 6b (Training Technologies).

In conducting the research described in this report, the investigators adhered to the policies and procedures set out in the Tri-Council Policy Statement: Ethical conduct for research involving humans, National Council on Ethics in Human Research, Ottawa, 1998 as issued jointly by the Canadian Institutes of Health Research, the Natural Sciences and Engineering Research Council of Canada and the Social Sciences and Humanities Research Council of Canada.

© Her Majesty the Queen as represented by the Minister of National Defence, 2006

© Sa Majesté la Reine, représentée par le ministre de la Défense nationale, 2006

Abstract

Dynamic Decision Making (DDM) is a skill that is both increasingly required and difficult to train for military commanders in today's security environment. Because it requires the timely sequencing of interdependent decisions in order to control complex and non-linear systems, DDM is a difficult skill to acquire for humans. Microworlds, which are stripped down simulations that focus on the dynamics of the target systems, have been proposed by many as training environments for DDM that avoid the time commitment, cost and personal danger of training command decision making with full-scale exercises or through mission experience. However, little research has been conducted on the factors that lead to effective microworld-based training. Specifically, it is unknown whether the time compression that occurs in microworlds enhances or inhibits the learning and transfer of complex system dynamics. A pilot study was conducted to examine whether participants are able to learn a simple DDM task in an accelerated microworld environment and then perform the same task in a similar but much slower environment. The results suggest that compressed-time microworlds can support training and transfer of DDM skills to "real-time" environments but that much remains to be learned about the conditions that favour the learning of DDM skills. Based on these results, general considerations for training DDM with microworlds and specific recommendations for improving the current study are provided.

Résumé

La prise de décision dynamique (PDD) constitue une aptitude très en demande mais difficile à enseigner pour les commandants militaires dans le contexte géopolitique actuel. La PDD est une aptitude difficile à acquérir, car elle requiert un séquençement rapide de décisions interdépendantes afin de contrôler des systèmes complexes et non linéaires. Les micro-mondes, des programmes de simulations simplifiées portant principalement sur la dynamique des systèmes cibles, ont été proposés par de nombreux milieux de formation de PDD afin de réduire l'investissement de temps, les coûts et le danger personnel associés à l'entraînement au moyen d'exercices militaires à grande échelle ou de l'apprentissage « sur le vif ». Cependant, les conditions qui favorisent l'efficacité de l'apprentissage au moyen de micro-mondes sont mal comprises. Plus précisément, on ne sait pas si la compression du temps dans les micro-mondes améliore ou limite l'apprentissage et le transfert des dynamiques de systèmes complexes. Une étude pilote a été menée afin d'examiner si les participants étaient capables d'apprendre une tâche simple de PDD dans un environnement micro-monde accéléré, puis accomplir cette même tâche dans un environnement semblable beaucoup plus ralenti. Les résultats de l'étude suggèrent que le temps compressé du micro-monde peut soutenir l'instruction et le transfert des aptitudes de PDD vers des environnements « en temps réel », mais il faut faire plus de recherches au sujet des conditions qui favorisent l'apprentissage de la PDD. D'après ces résultats, des observations d'ordre général sur la formation de la PDD et des recommandations précises pour de futures expériences sur ce sujet sont fournies.

This page intentionally left blank.

Executive summary

Accelerated Training for Command Dynamic Decision Making: A Pilot Study Using Microworlds

Jarmasz, J.; DRDC Toronto TR 2006-239; Defence R&D Canada – Toronto; November 2006.

Background: Dynamic Decision Making (DDM) is a skill that is necessary but difficult to train for command decision making in the Canadian Forces (CF). It is also a cognitively difficult skill, especially at the strategic and operational levels. Microworlds, which are stripped down computer simulations that focus on the dynamics of the target systems, have been proposed as training environments for DDM to mitigate the time commitment, cost and personal danger of on-the-job training or full-scale exercises. However, little research has been conducted on the effectiveness of microworlds as training tools. A pilot study was conducted to examine whether participants were able to learn a simple DDM task (resembling a simple peace support operation) in a microworld environment. The participants were separated into two groups, one where the task was trained in real time, and another where the task was trained in faster-than-real-time.

Results: The overall results show that all participants found the DDM task difficult. However, the accelerated training group performed as well or better than the real-time training group. The results suggest that compressed-time microworlds can support training and transfer of DDM skills for “real-time” environments, but much remains unknown about the conditions that favour the learning of DDM skills.

Significance: The pilot study suggests that dynamic pattern recognition can transfer across timescales. Were the results to be confirmed in a full-scale study, this would support the validity of microworlds for inexpensive and flexible faster-than-real-time training for command decision making in the CF.

Future plans: Short-term plans involve applying the lessons learned from this pilot study to a full-scale experiment on the transfer of DDM skills across time scales. Medium-term plans involve the identification of the application areas within the CF that would most benefit from microworld training, and the development of a research program to determine the factors and conditions that make microworld-based training for DDM effective. The long term goal is to develop a prototype microworld-based decision making training application that will model realistic CF strategic or operational decision-making scenarios.

This page intentionally left blank.

Sommaire

Formation accélérée de prise de décision dynamique pour le commandement : Étude pilote utilisant des micro-mondes

Jarmasz, J.; DRDC Toronto TR 2006-239; R & D pour la défense Canada – Toronto; November 2006.

Contexte : La prise de décision dynamique (PDD) est une aptitude nécessaire mais difficile à enseigner pour la prise de décision pour le commandement dans les FC. Il s'agit aussi d'une aptitude cognitive difficile, particulièrement aux niveaux stratégique et opérationnel. Les micro-mondes (des programmes de simulations simplifiées portant principalement sur la dynamique des systèmes cibles) ont été proposés comme environnements de formation pour la PDD afin de réduire l'investissement de temps, les coûts et le danger personnel de la formation au travail ou des exercices à grande échelle. Toutefois, l'efficacité des micro-mondes pour la formation est mal connue. Une étude pilote a été effectuée afin d'examiner si les participants étaient capables d'apprendre une tâche simple de PDD (semblable à une simple opération de soutien de la paix) dans un environnement micro-monde. Nous avons séparé les participants en deux groupes : le premier groupe a reçu une formation en temps réel et le deuxième a reçu une formation accélérée.

Résultats : Les résultats démontrent qu'en général les participants ont trouvé la tâche de PDD difficile. Néanmoins, le rendement du groupe qui a suivi la formation accélérée était égal ou supérieur à celui du groupe de la formation en temps réel. Les résultats semblent indiquer que la compression du temps qui a lieu dans les micro-mondes permet la formation et le transfert des aptitudes de PDD vers des environnements en temps réel, mais il faut encore pousser les recherches afin d'en apprendre plus sur les facteurs qui favorisent l'apprentissage des aptitudes de PDD.

Importance : L'étude pilote indique que la reconnaissance de motifs dynamiques peut se généraliser à des régimes plus lents. Une confirmation éventuelle de ces résultats dans une étude à grande échelle appuierait la validité des micro-mondes comme méthode de formation plus rapide, plus souple et moins coûteuse que la formation en temps réel de la prise de décision pour le commandement dans les FC.

Perspectives : Les plans à court terme visent l'application des leçons apprises lors de l'étude pilote dans une expérience à grande échelle portant sur la généralisation d'aptitudes de PDD à partir d'une formation accélérée. Il sera nécessaire, à moyen terme, d'identifier les domaines au sein des FC qui seraient les plus avantageés par une formation utilisant des micro-mondes et de développer un programme de recherche permettant de déterminer les facteurs et les conditions qui sont à la base d'une formation micro-monde efficace de la PDD. Enfin, l'objectif à long terme est de développer un prototype d'application pour la formation de prise de décision à l'aide d'un micro-monde qui présentera, de façon réaliste, des scénarios de PDD stratégiques et opérationnelles des FC.

This page intentionally left blank.

Table of contents

Abstract	i
Résumé	i
Executive summary	iii
Sommaire.....	v
Table of contents	vii
List of figures	ix
List of tables	x
Acknowledgements	xi
1. Introduction.....	1
2. Present Study	3
2.1 Objectives	3
2.2 Method.....	3
2.2.1 Participants.....	3
2.2.2 Decision-making task (scenario and model)	3
2.2.3 Apparatus and display	6
2.2.4 Procedure and design	7
2.2.5 Performance benchmarks	8
2.3 Results	13
2.3.1 Number of days stable.....	14
2.3.2 Resource losses.	16
2.3.3 Resource movements.	18
2.3.4 Individual patterns.....	20
2.4 Discussion	22
3. Implications and recommendations for training with microworlds	25
3.1 General considerations	25
3.2 Recommendations for improving the present study.	27
4. Conclusion	29
References	30
Annex A Microworld Models and Equation Listings	34
A.1 Model A1	34
A.1.1 Model Diagram.....	34
A.1.2 Equation listing.....	35
A.2 Model A2	36
A.2.1 Model diagram.....	36
A.2.2 Equation listing.....	37
A.3 Model B1	38

A.3.1 Model diagram.....	38
A.3.2 Equation listing.....	39
A.4 Model B2	40
A.4.1 Model diagram.....	40
A.4.2 Equation listing.....	41
List of symbols/abbreviations/acronyms/initialisms	43

List of figures

Figure 1: Sample microworld GUI.....	6
Figure 2: Benchmark run for model A1. The red horizontal line indicates the stability threshold.....	10
Figure 3: Benchmark run for model A2. The red horizontal line indicates the stability threshold.....	11
Figure 4: Benchmark run for model B1.....	12
Figure 5: Benchmark run for B2	13
Figure 6: Days stable by microworld model for the RT group (left panel) and the accelerated training group (right panel). Training data for model B1 were lost for the RT group. Training data for the transfer group is based on 2 participants only.....	15
Figure 7: Days stable for each model in the accelerated group training session. Due to lost data, only the averages of two participants for the last 3 training trials are shown.....	16
Figure 8: Resources lost by model for each training group. Training data for model B1 for the RT group and for one of the participants in the accelerated group are unavailable. ...	17
Figure 9: Resource losses for each model in the accelerated group training session. Due to lost data, only the averages of two participants for the last 3 training trials are shown.	18
Figure 10: Number of resource movements by model for each training group. Training data for model B1 for the RT group and for one of the participants in the accelerated group are unavailable	19
Figure 11: Number of movements for each model in the accelerated group training session. Due to lost data, only the averages of two participants for the last 3 training trials are shown.	20
Figure 12: Sample run by one participant from the transfer session showing a failure to recognize the destabilization caused by a prolonged security resource presence in Region A for model B2 (one of the “stabilize-and-retreat” models). The horizontal red line at 15 units marks the desired stability level.	21
Figure 13: Sample run by one participant from the transfer session showing oscillations in the stability index due to backlash from excessive resource movements (i.e., a drop in resource levels without a corresponding withdrawal of resources by the participant) for model A1, a “stabilize-and-stay” model requiring slow resource build-up. The horizontal red line at 15 units marks the desired stability level.....	22

List of tables

Table 1: The four microworld models used and their associated resource policies	5
Table 2: Results from benchmark runs	9

Acknowledgements

The author would like to acknowledge the help of HumanSystems Inc. (Guelph, Ontario), and in particular Ms. Lora Bruyn, in running the pilot study, collecting the data, and compiling the participants' subjective impressions of the study. The author also thanks Mr. Rick Bodner, Ms. Kim Wulterkens and Ms. Marie-Ève Jobidon for reviewing drafts of this report.

This page intentionally left blank.

1. Introduction

The complex and dynamic nature of the operations-other-than-war (OOW) (e.g., peace support, the 3-block war concept) in which Canada and allied nations are increasingly involved requires commanders in the Canadian Forces (CF) to call upon dynamic decision making (DDM) skills to an unprecedented degree. As noted by Brehmer [1], [2], Clancy et al. [3], and others, dynamic decision making tasks are usually characterized as tasks that require a series of interdependent decisions, whose state change both autonomously and as a result of the decision maker's actions, and where decisions have to be made in real time. DDM is a skill that is notoriously difficult for human beings, even after many years of experience (Brehmer [1], Dörner [4], Sterman [5]). As Bakken and Vamraak [6] have shown, DDM is a challenge even for experienced military analysts. There is therefore a need to develop effective training programs for DDM in the CF.

It is difficult to develop training opportunities to practice DDM for military commanders, especially at the level of strategic and operational decision making. Bakken & Gilljam [7] note that commanders typically learn strategic and operational DDM through experience in joint or multinational operations, but for many smaller nations (e.g., Canada, or the authors' native Norway) such campaign experience tends to be rare; furthermore, the full-scale campaign exercises that are required to train this type of DDM are extremely expensive and time consuming. In a recent study of the applicability of systems dynamics concepts to staff officer training in the CF, Rehak et al. [8] note that training at the Canadian Forces College includes some concepts relevant to DDM (e.g., some academic discussion of second- and third-order effects) but does not include a systematic approach to teaching the non-linear, holistic "systems thinking" skills that are thought to underlie proficiency in DDM.

One promising solution to the problem of training DDM skills is the use of microworlds. These are simulated interactive models that capture the high-level dynamics of relevant DDM situations while stripping away unnecessary detail (Brehmer [1], Haberstroh et al. [9], Senge [10]; Shanteau et al. [11], Sterman [5]). Microworlds allow people to experience the dynamics underlying a complex DDM situation within a compressed timeframe. It is hypothesized that the time compression allows people to rehearse DDM skills more often and makes it easier to learn correct cause-and-effect relationships despite the feedback delays inherent in the systems that are modeled (Bakken & Gilljam [7]). Thus microworlds could provide an inexpensive, repeatable, adaptable, and pedagogically effective surrogate for experiential learning or "learning by doing" for military command decision making.

Microworlds have been extensively used to study DDM (see Gonzalez et al. [12] and Rouwette et al. [13] for reviews). However, while there is some evidence that microworlds can improve DDM skills, there has not been much investigation of microworlds as training tools. Little is known about how to design microworlds to reliably produce effective learning (Gonzalez [14], Gonzalez & Quesada [15], Lerch & Harter [16]). Furthermore, little research has been done on whether learning in microworlds transfers to real-life situations, which are both more complex and unfold at a much slower pace than microworld simulations. While there is some research on accelerating time in high-fidelity simulator-based training (e.g., Ali et al. [17] studied time compression to improve combat flight training), Levy et al. [18] note that most reports on training with microworlds consist of anecdotal reports from the domain of business management. To date, experimental research on transfer of training in microworlds appears to have been limited to a few studies in the realm of primary and secondary education (Martin et al. [19], Miller et al. [20]). Research examining the use of microworlds for military training (Bakken & Gilliam [7],

Bondanella et al. [21], Levy et al. [18]) has not yet progressed to the point of conducting transfer-of-training studies.

Research suggests that DDM relies heavily on recognition processes (Gonzalez et al. [22], Gonzalez & Quesada [15]), as do most types of expert naturalistic decision making (Klein [23], [24]). The recognition of temporal patterns suffers when patterns are presented at speeds different than those at which they were learned (Boltz [25], Schulkind [26]), but it is not known how important this effect is. If the ability to recognize dynamic patterns and make use of them for DDM requires those patterns to have been learned at or near to their real-time speed, then microworld-based training will be of limited use; at best it will serve to train those aspects of decision making that not dependent on temporal dynamics, if such aspects can be identified for DDM. Therefore, it is of paramount importance for the future of microworlds as training tools to determine the extent to which DDM skills that are learned at an accelerated tempo can be applied in other dynamic environments that unfold at a much slower pace.

2. Present Study

2.1 Objectives

The objective of this study was to pilot an experimental design and procedures in preparation for a series of experiments aimed at determining whether DDM skills can be learned in an accelerated training context and then transferred to a real-time environment. To this end a set of related microworlds that require the identification of temporal patterns in order to determine the correct decision strategy were developed (see Section 2.2.2). However, given the well-documented problems humans have in performing DDM tasks, and in particular the fact that improvement in DDM performance appears to be difficult to induce, it was not clear whether participants would be able to master the microworlds designed for the experiment to a degree sufficient for transfer effects to be detectable. A second problem was that running a “real time” transfer session with a duration similar to that of strategic or operational command decision making (i.e., many days or even weeks) presented logistical difficulties that were not worth addressing until questions concerning the experimental procedures had been resolved. Therefore, a pilot study was run to determine whether participants would be able to learn and perform the DDM task under the training regimen designed for the experiment, and then perform the task in a microworld environment that was slower than the training environment but faster than typical operational and strategic decision-making.

2.2 Method

2.2.1 Participants

Six volunteers were recruited from the general population in Guelph, Ontario by a contractor (HumanSystems Inc.). The sample included 3 men and 3 women (mean age 35.2 years, range from 21 to 57 years). Participants were recruited and compensated according to DRDC guidelines for human participation in experimental research (DRDC Toronto [27], Pigeau [28]). Three of the participants had previous experience with strategy/simulation computer games.

Due to the sustained nature of the cognitive effort required in this experiment, participants were asked to ensure they were well rested for the duration of the experiment, and not to modify their usual consumption of caffeine, tobacco or prescription medicines for the duration of the experiment. They were requested not to consume substances affecting mental capacities (e.g., alcohol, recreational drugs) 48 hours prior to and for the duration of the experiment, and to inform the investigators should a medical condition arise during the experiment that might affect their mental performance.

2.2.2 Decision-making task (scenario and model)

Participants were asked to perform a simple dynamic decision-making task inspired by the peace support and reconstruction scenarios that are the focus of many current CF operations. In general terms, participants were asked to build and maintain a security force over a period of 91 simulated days in a notional war-torn region, called “Region A.” The goal of each trial was to stabilize and reconstruct the region, where unnamed insurgents were known to be operating. The Stability Index (an arbitrary numerical representation of the stability of Region A) varied as a function of

DRDC Toronto TR 2006-239

the strategies that were used to move security force resources in and out of the region. Participants were asked to manage security resources so as to maintain the security index at or above a criterion value of 15 units. Details of the relationship between security resource movements and the stability index are discussed below. Only high-level aspects of the mission were represented: for instance, no distinction was made between security forces and their equipment, or between different methods of stabilizing the region (e.g., humanitarian actions vs. providing security and policing). Participants adjusted the stability index by moving discrete amounts of security resources in or out of the region at times of their choosing. There was no representation of different ways of applying the security resources (e.g., humanitarian actions vs. security/policing activities). The only information presented to the participant was: time elapsed; amount of resources allocated to them (notionally located at “base camp”) at the start of the trial; how many security resources were present in the affected region at a given time; and the stability index of the region at a given time. Time in the simulated world progressed continuously, and the different values presented to or controlled by the participant were updated at the start of every day in the simulation.

The main hypothesis motivating this study is that people perform DDM tasks relying heavily on pattern recognition processes (Gonzalez et al. [22], Gonzalez & Quesada [15]), and that it is important to determine whether participants are able to learn dynamic patterns at one speed and then apply this training in support of DDM at a much slower speed. The experiment was therefore designed so that participants needed to learn to identify certain dynamic patterns in the microworld and apply a decision strategy appropriate to the patterns. To that end four different dynamic models, each producing specific dynamic patterns, were created to drive the microworld task described above. Each model resulted from the combination of one of two resource movement policies and one of two policies for maintaining resources in Region A. The two resource movement policies were: a “slow” policy where movements of 3 or more units at a time results in a sharp decline of the resources already in theatre after a delay of a few days (representing local resentment and backlash generated by a sudden foreign presence), and a “fast” policy where movements of 3 units or less at a time resulted in units lost in transit (representing the concept of strength in numbers). The two resource maintenance policies were: a “stabilize-and-stay” policy where the stability index is proportional to the number of security resource units in Region A and maintaining stability therefore requires maintaining a minimum amount of resources in the region throughout the mission, and a “stabilize-and-retreat” policy where a critical mass of security resources triggers a rise of the stability index to a set level (modeled as the activation of a regulatory feedback loop) but where a continued presence of the security resources eventually produces animosity and retaliation from insurgents (which overwhelms the regulatory feedback loop). This requires the security resources to be withdrawn before the insurgents’ animosity rises, but the effects of the initial stabilization effort eventually wear off (i.e., with time the regulatory feedback loop becomes deactivated) and thus security resources need to be moved in and withdrawn periodically over the course of the mission.

These different policies are intended to illustrate the notion that the way in which a security force establishes itself in a region (taking time to build a relationship with the locals vs. moving quickly against insurgents) can have different consequences depending on the characteristics of the region in question. The policies for resource movement reflect the idea that in some cases moving too fast will draw unwanted attention from insurgents or generate ill-will from the locals, and negatively affect the mission (increase casualties or destabilize the region), whereas under other circumstances moving too gradually will expose the security force to unnecessary risk. The policies for resource maintenance reflect the contrast between dependent regions where the local

population cannot maintain stability in the region without help from the security force (a withdrawal leads to collapse) and more autonomous regions where a continued presence initially increases stability but with time generates resentment from the population and leads to destabilization of the region.

The four models generated from the combination of the resource policies (see Table 1) were labelled A1 (slow movements, stabilize-and-stay), A2 (slow movements, stabilize-and-retreat), B1 (fast movements, stabilize-and-stay), B2 (fast movements, stabilize-and-retreat). The details of the models are given in Annex A.

Table 1: The four microworld models used and their associated resource policies

Resource maintenance policy	Resource movement policy	
	Move slow or risk backlash	Move fast or risk attacks in transit
Stabilize-and-stay	Model A1	Model B1
Stabilize-and-retreat	Model A2	Model B2

During the training session the participants' task was to learn to apply the appropriate decision strategies (i.e., the resource management policies) to each microworld model, and in doing so to implicitly learn to recognize the specific dynamics of each microworld model. To this end, participants were given instructions for each model that suggested the correct decision strategy without prescribing it (a sample of the instructions is shown in Figure 1). The intent was to point participants in the right direction while still encouraging exploratory behaviour. Research has shown that people are unlikely to learn the dynamics of a complex system merely through trial-and-error along with performance feedback, and that some type of indication of "correct" or preferred behaviour is helpful (Gonzalez [14]). However, the dynamics typical of each model are best learned through exploratory behaviour. For instance, the resource movement policy governing a particular model can be determined by a participant by making deliberately large resource movements and observing whether a backlash occurs after a suitable delay, and the resource maintenance policy can be determined by observing whether the stability index tracks the level of security resources Region A (with a delay of a few days). Thus it was hoped that by providing vague and suggestive instructions rather than detailed and prescriptive ones, participants would learn to apply the appropriate policies while still engaging in exploratory behaviour in order to learn to recognize the dynamics of the different models.

In the transfer session participants were required to interact with microworlds based on the same four models as in the training sessions, but the same generic instructions, recapitulating all the different resource policy suggestions from the training sessions, were given regardless of the model used. Participants were thus unable to identify the model from the instructions and had to

rely on their memory of the dynamic patterns learned during training in order to select and apply the correct decision strategies.

2.2.3 Apparatus and display

The microworld was constructed by the author and presented to the participants using commercial software for creating interactive simulations (iThink version 8, from isee systems Inc.¹). This system dynamics modeling application was used both to construct the four models discussed above as stock-and-flow systems, and also to construct the graphical user interface (GUI) that the participants used to manipulate the microworld. The models (stock-and-flow diagrams as well as equation listings) are shown in Annex A. Figure 1 illustrates the appearance of the GUI.

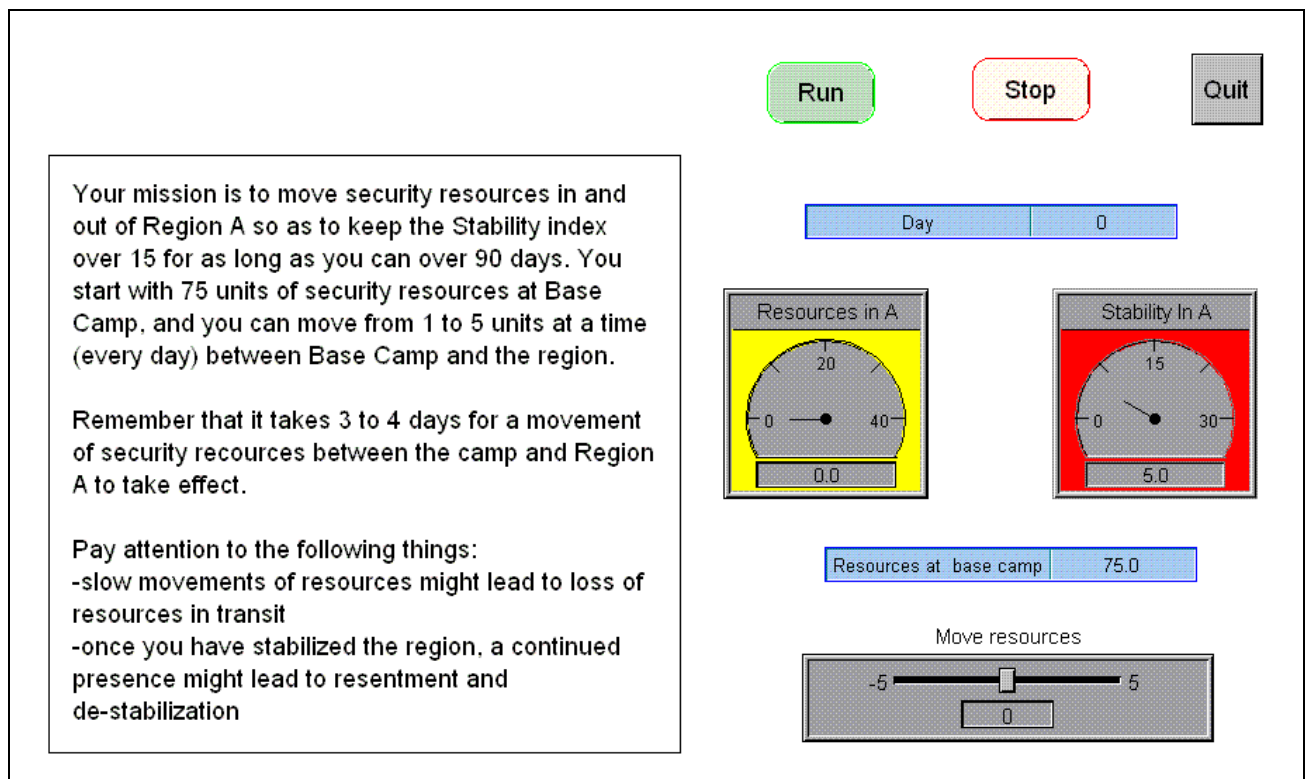


Figure 1: Sample microworld GUI

The left side of the GUI provided a brief cover story and the suggestions for resource management appropriate to the model underlying the microworld (as noted above, in the transfer session the GUI showed generic suggestions regardless of model used). The right side of the GUI had two needle gauges that displayed the states of two variables of the microworld, resource level

¹ A newer version of iThink (version 9) was available to the author at the time of the study; however, it was found that bugs in the later version of iThink prevented accurate updating of the model variables displayed on the GUI, and the earlier version of iThink was retained for the experiment.

in theatre and stability index of the troubled region. Two numeric read-outs displayed the number of simulated days elapsed and the number of resources left at the notional base camp. The intention was to encourage participants to rely as much as possible on their internal sense of timing and rhythm in performing the task, while keeping the number of dynamic variables they need to track to a bare minimum. Therefore, a gauge format rather than a continually updated graph of the variables over time (also available in iThink) was chosen because it conveyed the dynamics of the variables without providing a record of the evolution of the microworld over time. Such a representation might have encouraged the participants to develop speed-invariant mental models of the microworld dynamics, which might have interfered with their reliance on their sense of temporal dynamics. As discussed below, the gauge format was probably not optimal for learning the DDM skills, and future research will need to focus on optimal representation of system dynamics in support of DDM training and transfer.

The right side of the GUI also provided the participants with a slider mechanism (actuated via a computer mouse) for deciding how many resource units to move each day between the base camp and Region A. The slider would reset to zero at the beginning of each simulated day. Participants were allowed to control only one variable in able to keep the experimental task relatively simple, but future research will have to provide more realistic microworlds where many variables (representing different types of decisions) can be controlled.

2.2.4 Procedure and design

Participants were asked for their written consent before participating in the study. The reconstruction scenario was presented in general terms and the principles underlying the microworlds were explained. They were then asked to practice applying these principles by interacting with all four microworlds in a training session. As described above, the instructions displayed on the computer screen informed the participants as to the appropriate decision strategies for each training trail, and they were verbally encouraged to explore for themselves the effects of different actions in each microworld model, while at the same time trying to keep the stability index above the criterion value for as long as possible.

Participants were randomly assigned to one of two training groups, defined by the speed at which the microworld is presented during the training trials. The first group (called the “accelerated training” group) interacted with the microworld presented at a speed of one simulated day lasting 3 seconds. Thus, each training trial lasted 270 seconds (or 4.5 minutes). This represents an acceleration factor of 20 relative to the transfer session (described below), where each day lasts 60 seconds. In this group, participants had the opportunity to interact with each microworld four times during the two hours that the training session lasted, for a total of 16 training trials.

The second group (designated the “real-time group” or “RT group”) trained with the microworlds presented at the same speed as in the transfer session. Thus each training trial for the RT group lasted 90 minutes, and only 4 practice trials (one per model) were given to participants in the RT group. This manipulation was meant to achieve two things: the RT group’s training trials were meant to represent performance of the DDM task in the absence of training, and the RT group’s transfer trials were meant to represent the effects of training in real time (contrasted with training in accelerated time).

Following the training session, on a separate day, participants were asked to participate in a second session (the “transfer session”) where they performed the DDM task once with each of the four models (for a total of 4 transfer session trials), presented in random order. They were not

informed as to which specific model they were interacting with and had to rely on the dynamic patterns they learned in the training session to recognize the model and apply the correct decision strategies in order to keep Region A stable for as long as possible. The speed of the simulation during the transfer session (one day of simulated time for every 60 seconds of real time) served as the reference speed (i.e., “real time”) for the training sessions.

The experimental design involves one independent measures factor (training group), and two repeated measures factors: session (training and transfer) and microworld model (A1, A2, B1 and B2). This last factor could be further divided into 2 factors (resource movement policy and resource maintenance policy), but due to the low number of participants in the study model will be treated as a single factor. As the accelerated training group experienced four training trials for each model (versus one in the RT group), the fourth trial for each model was retained when comparing the training and transfer sessions.

According to a strict recognition-primed decision hypothesis, participants who train the DDM task in the RT condition should perform better in the transfer session than those who trained in the accelerated condition. Nevertheless, the hypothesis that accelerated training is possible at all predicts that the accelerated group’s performance in the transfer session will be better than the training session performance of the RT group. According to the theory that accelerated training has an advantage over RT training (either because of extra repetitions or because the compressed time scale allows the dynamics of the system to be perceived more easily²), the transfer session performance of the accelerated group should be better than the transfer session performance of the RT group.

2.2.5 Performance benchmarks

There was no way of knowing *a priori* what level of performance could be reasonably expected from participants in the study, and what measures would be useful in assessing participants’ performances. In order to determine a likely upper bound on performance and to investigate useful measures, benchmark runs were performed by the author attempting to achieve the best possible performance with full knowledge of the microworld’s model’s dynamics (therefore as close as possible to an optimal or ideal performance level), and analyzed to identify useful measures.

Figures 2 to 5 show the stability level, resource movements and actual resource levels in Region A over time for the benchmark runs for each microworld model. The figures also display the stability threshold (arbitrarily set to 15 units) above which the region was deemed to be “stabilized.” From the figures, the number of days that Region A was kept stable (i.e., that stability was above 15 units) was the most obvious performance measure. A visual inspection of Figures 2 to 5 suggests that under ideal conditions participants should be able to keep the region stable for a high proportion of the simulated run for all microworlds (from a minimum of 67 days for model A2 to a maximum of 82 days for model B1) with relatively low losses (approximately 5 units lost over the course of one simulation run for most models except for B1³). Resource losses

² The potential effects of increased number of repetitions and better dynamics perception in the accelerated training regimen are confounded in this design, and will have to be examined separately in future research.

³ Model B1 was accidentally left with a higher rate of baseline resource attrition from earlier iterations of the model. As the error was not deemed to qualitatively affect the performance of the DDM task, it was not corrected for the pilot study.

were selected as one measure of the appropriateness of the resource movement policies chosen by participants (inappropriate resource movements should result in higher losses than in the benchmark runs). The main distinction between the models is in the resource management strategies that need to be applied to keep Region A stable. For the “stabilize-and-retreat” models (A2 and B2), optimal resource movements take on a periodic, saw-tooth shaped pattern, with the period of the pattern being longer in model A1 (due to that model’s requirement that resources be moved slowly to avoid an insurgent backlash). Generally speaking, the “stabilize-and-retreat” models require more movements than the “stabilize-and-stay” models, and the slow movement models (A1 and A2) require smaller and therefore more frequent movements as well; these differences between benchmark runs were captured by computing the total number of resource movements (regardless of movement size or direction).

The values of the three measures (number of days stable, cumulative losses, cumulative resource movements) for the benchmark runs are summarized in Table 2. These were the main measures selected for the study. Note that measures commonly used in control-theory type studies (e.g., the root mean square error between system performance and the reference signal) were not used because the “reference signal” in this case, the stability criterion of 15 units, did not vary throughout the experiment, and there was no direct incentive for participants to keep the stability index below the criterion value; that is, there was no need to evaluate how well participants tracked a reference signal.

Table 2: Results from benchmark runs

Measure	Model			
	A1	A2	B1	B2
Time stable (days)	70	67	82	78
Losses (arbitrary units)	5.22	5.52	15.15	5.8
Number of resource movements	23	76	7	49




Figure 2: Benchmark run for model A1. The red horizontal line indicates the stability threshold.




Figure 3: Benchmark run for model A2. The red horizontal line indicates the stability threshold.




Figure 4: Benchmark run for model B1




Figure 5: Benchmark run for B2

2.3 Results

Data for the three main measures (days stable, resource losses, number of resource movements) were computed for each participant for the different conditions in the experiment. Due to the small number of participants in the pilot study, only descriptive statistics (means) and qualitative patterns will be reported, with the caveat that these do not carry the inferential weight that the usual ANOVA would. Also, due to technical difficulties, some data were lost from some of the training trials (the trials for model B1 for two participants in the RT group, and the first training trial for all models for 2 participants in the accelerated group as well as all training data for one participant in that group). This further complicates the interpretation of the results. It also prevented one of the comparisons originally planned for the study, that is, the comparison of the transfer session performance for the accelerated group and the training session performance for the RT group. Fortunately the accelerated group's transfer session performance compared favourably with the RT group's performance in the same condition, and evidence for skill transfer

in the accelerated group did not require a comparison with the RT group's training session performance.

2.3.1 Number of days stable

The primary performance measure was the number of days (out of 91⁴) that the participants managed to keep Region A stable. An overall comparison (ignoring specific microworlds) of performance in the transfer session suggests that performance was better in the accelerated group (41.2 days stable vs. 26.3). An examination of performance by microworlds shows that the accelerated group outperformed the RT group mainly on models B1 (77.3 vs 21.6 days stable) and A1 (56.6 vs 41.3), whereas the RT group scored better for model A2 (33 vs 17.3) and both training groups had similar scores for model B2 (13.3 for the accelerated training group vs 9.3 for the real-time group. Note that the only condition where performance came close to matching the benchmark was model B1 for the accelerated group (benchmark = 82 days). This is a first indication that the accelerated group learned something that the RT group did not, in particular for the “stabilize-and-stay” models (A1 and B1).

The days stable measure was examined for transfer effects for both training groups, but because of missing data, transfer for model B1 could not be examined for the RT group, and training data were available for only 2 of the 3 participants in the accelerated training group. Performance for the training and transfer sessions are shown in Figure 6. Excluding model B1, overall performance (averaged over microworlds) was low (below 30 days stable) and did not differ much between training groups, though it appeared to improve slightly for both groups in the transfer session. Thus, the effect of transitioning from the training to the transfer task was about the same in both groups, regardless of training speed. With regard to individual microworlds, the performance patterns show little variation between the training and transfer sessions for each group. The slight improvement seen for model B1 for the accelerated group is consistent with the overall pattern for the other training group and models.

⁴ Nominally the simulations lasted 90 days but since iThink was set up to start counting time from Day 0, 91 days were used to end the simulations on Day 90 and to beautify output graphs. Future versions of the microworld will start counting time from Day 1.




Figure 6: Days stable by microworld model for the RT group (left panel) and the accelerated training group (right panel). Training data for model B1 were lost for the RT group. Training data for the transfer group is based on 2 participants only.

Performance was also examined across training trials in the training session for the accelerated training group (shown in Figure 7). As stated, training data were available for only 2 participants in the accelerated training group, and only the last 3 learning trials (2 to 4) were available for each participant; the data therefore do not give a complete picture of what happened in the training trials, especially since the crucial first trial was lost. Overall the data suggest some improvement from trial to trial when all performance is averaged over all the microworlds (number of days stable goes from 27.8 on trials 2 and 3 to 36.6 on trial 4). Examination of stability performance according to microworld and trial suggests that performance is uneven across models and across trials. Performance is highest for models A1 and B1 (stabilize-and-stay), reflecting their relatively simple resource management policies. However, only B1 shows consistent improvement from trial to trial (35.5 to 74 to 75 days, nearing the benchmark of 82), whereas performance in A1 drops on the third trial and recovers partly in the last trial (57.5, 27.5, 51 days respectively). Models A2 and B2 show similar patterns of drops in performance in trial 3, with either a partial recovery (B2) or an improvement over trial 2 (A2).




Figure 7: Days stable for each model in the accelerated group training session. Due to lost data, only the averages of two participants for the last 3 training trials are shown.

2.3.2 Resource losses.

The mean number of resource units lost in the transfer session was higher for the accelerated group (20.7 units) than for the RT group (16.9 units). Examining losses for the accelerated group for the different microworlds, there seems to be an observable pattern for the accelerated group (see Figure 8, right panel). The pattern consisted of high losses for the slow movement models (37.7 units for A1, 30.6 for A2), exceeding in both cases the benchmark values, and low losses for the fast resource movement models (11.3 units for B1 and 3.2 for B2, unexpectedly below the benchmark values). In the RT group (Figure 8, left panel), losses exceeded the benchmark for all models except B2 and were highest for the “stabilize-and-stay” models (20.5 units for A1 and 28.6 for B1).

Resource losses were compared for transfer effects for both groups, again bearing in mind the missing data as reported above for the days stable measure. The RT group showed a reduction in resource losses between sessions (20.8 units in the training session vs. 16.91 in the transfer session), with most of the reduction due to a drop in losses for model A1 (43.5 to 20.5 units; see Figure 8, left panel). The accelerated training group showed a slight increase in resources lost

between sessions (17.7 units in the training session vs. 20.7 in the transfer session), due mostly to more losses for models A1 and A2 in the transfer session (see Figure 8, right panel).




Figure 8: Resources lost by model for each training group. Training data for model B1 for the RT group and for one of the participants in the accelerated group are unavailable.

Resource losses were also examined across training runs for the accelerated group, and are shown in Figure 9. Losses (averaged over all across microworlds) were above the benchmark average for all models (7.9 units lost) and did not exhibit appreciable improvement (i.e. no noticeably drop) during the training session (31.8, 33.6 and 29.8 units lost for runs 2, 3 and 4 respectively). Improvements were seen only for models A1 (51.8, 51.4, and 44.3 units lost) and B1 (32.7, 19.9, 18.9 units lost). Resource losses increased during the training session for model A2 and fluctuated (up then down) for model B2.

In principle, the trade-off between stability performance (number of days stable) and resource losses could be examined to assess the efficiency of participant performance. For instance, two participants might have achieved similar stability performances while incurring different resource losses, with lower losses indicating more efficient performance. This could be assessed by computing the ratio of days stable to resource losses for each participant, with a higher ratio indicating more efficient use of resources in the stabilization task. This measure was computed

DRDC Toronto TR 2006-239 17

for the participants, but it did not show any patterns or differences in addition to those already observed in the number of days stable and resource loss measurements. Furthermore the ratios were difficult to interpret without referring to the other two measurements, given the wide variability across conditions in these measurements. The ratios are therefore not reported here. However, this measure will be considered in future research with this experimental platform.




Figure 9: Resource losses for each model in the accelerated group training session. Due to lost data, only the averages of two participants for the last 3 training trials are shown.

2.3.3 Resource movements.

The number of resource movements were computed for each condition where data were available and are shown in Figure 10. Average number of movements during the transfer session were higher for the accelerated group than for the RT group (61.2 movements vs. 24). For the accelerated group in the transfer session, number of movements were higher for the stabilize-and-retreat models (79.7 for model A2 and 76.7 for B2) than for the stabilize-and-stay models (56.3 movements for A1 and 32 for B1), thus approximating the qualitative pattern set in the benchmark runs. Note however that the pattern seen in the benchmarks of lower resource

movements for the fast resource movement models (A2 and B2) is not apparent for the accelerated group. The RT group, however, showed little variability from model to model.

Number of resource movements were also examined for transfer effects for both groups. The average number of movements increased in the transfer session for the accelerated group (59 movements vs. 49.1 in the training session), and it decreased for the RT group (24.1 movements in the transfer session vs. 39.4 in the training session). The direction of change between training and transfer session was consistent across models for both groups.




Figure 10: Number of resource movements by model for each training group. Training data for model B1 for the RT group and for one of the participants in the accelerated group are unavailable

Resource movements were also examined across training runs for the accelerated training group (Figure 11). Only models A1 and B1 showed consistent movement towards the bench mark value, whereas models A2 and B2 showed oscillations similar to the ones seen above for the days stable and resource losses measures.




Figure 11: Number of movements for each model in the accelerated group training session. Due to lost data, only the averages of two participants for the last 3 training trials are shown.

2.3.4 Individual patterns

Since only 6 participants were involved in the study, their individual performance logs were examined for any interesting quantitative patterns that were not captured by the measures discussed above.

Some recurring patterns of note were identified in the individual performance logs. Here we consider only the transfer session. One pattern occurred in the “stabilize-and-retreat” models (A2 and B2) and involved participants persisting in maintaining an inappropriate strategy and failing to move out resources after stability is achieved, and even increasing resources in theatre as stability falls (shown in Figure 12 for a participant in the RT group). Such persistence is commonly seen in DDM research and complex problem solving (Brehmer [1], [2], Dörner [4], Serman [5]). Another pattern of note occurred when participants in the “stabilize-and-stay” models produced oscillating patterns instead of keeping the stability index constant (shown in Figure 13 for a participant in the accelerated training group). In some cases, as in Figure 13, this was due to participants not realizing right away they were moving too many resources at once and causing a backlash (i.e., the reduction in resources was not directly decided by the participant). In

others, participants appear to have moved resources out of Region A even though it was stable, perhaps because they thought that they were interacting with one of the “stabilize-and-retreat” models. Whatever the reasons for the oscillations, it is known that participant-induced oscillations are common when people interact with systems that have long delays (Brehmer [1], [2], Dörner [4], Jagacinski & Flach [29]).




Figure 12: Sample run from the transfer session by one participant (RT group) showing a failure to recognize the destabilization caused by a prolonged security resource presence in Region A for model B2 (one of the “stabilize-and-retreat” models). The horizontal red line at 15 units marks the desired stability level.

Another noteworthy pattern appeared when comparing participants from the two training groups. Participants in the RT group displayed a tendency to use large resource movements (typically the largest possible movement, i.e., 5 units at a time) regardless of whether the model they were interacting with required a slow or a fast resource build-up, as seen in Figure 12. The participants in the accelerated training group used a wider range of resource sizes and thus showed less consistent use of large resource movements, as is shown in Figure 13. To verify this observation, the average number of resource units per movement was calculated for both training groups. The RT group had an average of 4.0 units per movement ($SD = .75$), whereas the accelerated group had an average of 2.5 units per movement ($SD = .52$), confirming the difference observed in the individual runs. Furthermore, these values varied little as a function of the model underlying the

microworld, suggesting that participants did not adjust their preferred movement size as a function of the resource movement policy required by the model.

Turning to individual differences, it is worth noting that while the accelerated group outperformed the RT group, and did so in particular for model B1, one participant in the RT group (the only one in the RT group with previous experience with simulation-type computer games) performed on the whole as well if not better than the accelerated group. In fact, that participant outperformed all the accelerated group participants for model A2, one of the “stabilize-and-retreat” models that were the most difficult to control due to the backlash process.




Figure 13: Sample run from the transfer session by one participant(accelerated training group) showing oscillations in the stability index due to backlash from excessive resource movements (i.e., a drop in resource levels without a corresponding withdrawal of resources by the participant) for model A1, a “stabilize-and-stay” model requiring slow resource build-up. The horizontal red line at 15 units marks the desired stability level.

2.4 Discussion

On the whole, participants did not achieve the performance benchmarks, except for participants in the accelerated group interacting with model B1, which is perhaps the simplest model and thus easiest to learn. It is worth noting that resource losses and resource movements in general tended to be quite a bit higher than the benchmark values. This underlines both the difficulties (discussed above) that people typically encounter in trying to manipulate complex systems with delayed

feedback loops and the inadequacy of the training regimen in this pilot study. Nevertheless, the fact that the accelerated training group performed as well and sometimes better than the RT group suggests that transfer from an accelerated training regimen to a real-time transfer task is possible. Given the generally poor performance in the study, and the fact that the performance measures did not seem to vary much between the transfer and the training sessions, it is reasonable to ask whether participants learned anything during the training session. If so, did the accelerated group learn to perform the task differently? However, it might rather be the case that participants did not learn anything and that their performance merely reflects their “naïve” understanding of the task, in which case the differences between the groups merely reflects sampling error.

This issue cannot be addressed satisfactorily without running a study with an adequate number of participants. Nevertheless there are signs that the participants did engage in some form of learning, and that the accelerated training group did in fact learn decision-making strategies that the RT group did not. The first sign is that the accelerated group consistently outperformed the RT group on model B1 in the transfer session. The second is that the accelerated group showed consistent improvement for the days stable and resource losses measures for model B1. The third is that both groups showed improvement on the days stable measure for the transfer session relative to the training session, suggesting that learning was occurring and that it was continuing to occur in the transfer session. The fourth is that the accelerated group produced more consecutive movements in the same direction (i.e., more organized sequences of movements), in particular for the “stabilize-and-retreat” models, than did the RT group, suggesting that while they did not master the appropriate resource movement strategies they were attempting to apply them. Fifth is the fact that each group consistently used different sized movements, even in the transfer session, indicating that different training speeds can produce different decision strategies that might survive into the real-time environment. Taken together, these observations suggest that while the training was inadequate for participants to approach the performance benchmarks, some learning did take place in the microworlds, and the accelerated training group was able to learn to do some things better than the RT group did. The results also suggest that some forms of negative transfer might occur with microworld training. For instance, the different movement sizes that training groups seem to have adopted regardless of model suggests that the particular simulation speed used in training might encourage the learning of decision strategies that are inappropriate for the target environment.

As discussed above, a review of the individual performance logs revealed two recurring behaviour patterns that were “maladaptive” in that they contributed to poor performance: participant-induced oscillations in the “stabilize-and-stay” models and increases in resources despite declining stability in the “stabilize-and-retreat” models (i.e., perseverance in applying an inappropriate behaviour). As noted above, these patterns have been extensively documented in studies of DDM and complex problem solving with microworlds. Thus even with the very simple microworlds used in this study it was possible to induce the classic difficulties that humans exhibit during DDM. This suggests that interventions to improve DDM should not underestimate the difficulty of dynamic decision making tasks.

Another noteworthy finding that was not captured by the three main measures in the study was the difference in the average movement size between the training groups, and the fact that the size of movements did not seem to vary much in the transfer session as a function of the model underlying the microworld. The large movements used in the RT group are suggestive of some non-proportional control regimes that are commonly seen in control systems (either human-controlled or automated) with long feedback delays. As discussed in Jagacinski and Flach [29], the control signal (which corresponds to the decision behaviour in a DDM task) in typical

feedback-control systems with short delays is generally proportional to the error signal (i.e., the difference between the intended state and the actual state of the system); this is known as proportional control. When feedback delays become long, proportional control generally results in oscillations and system instability, because the error signal available to the controller reflects past states of the system, and the controller thus “reacts” too late. In such cases some kind of non-proportional control is required. One specific type of non-proportional control strategy, called “bang-bang” control, involves making discrete, all-or-nothing adjustments to the system, observing the effects, and making further discrete adjustments followed by observation until the desired state is achieved. The large movement size used by participants in the RT group suggests they might have adopted a “bang-bang” control strategy. Another type of non-proportional control involves making continuous adjustments to the system that are based on an internal model of the system, allowing the controller to anticipate future states and take into account feedback delay rather than reacting to the current error signal (which in fact reflects the system’s past). The smaller resource movements used by the accelerated group participants suggests that they attempted to use either proportional control or non-proportional control based on an internal model (with only a limited degree of success). The fact that a difference in movement size was observed between the groups suggests that the real-time duration of the feedback delays (rather than their duration in simulated time) determine whether human controllers will adopt continuous or bang-bang control strategies. The fact that they were observed in the transfer session suggests that control strategies transfer to post-training contexts and might enhance post-training performance (positive transfer) or interfere with it (negative transfer) depending on the context.

3. Implications and recommendations for training with microworlds

This pilot study had the objective of determining whether participants were able to learn to recognize different dynamic patterns, and use these to select appropriate decision strategies for a simple DDM task. It also aimed to determine measures for evaluating the transfer of these DDM skills from an accelerated training environment to a slower transfer environment. The training microworld and the training regimen were designed to encourage reliance on the temporal dynamics of the microworld, with a view to future experiments designed to examine whether DDM skills can transfer from an accelerated microworld environment to a real-time one. Thus, the study was not designed specifically to optimize learning in the training session. Unsurprisingly, therefore, participant performance in the study was suboptimal. An examination of the results suggests some general considerations for optimizing training with microworlds, which appear to be in short supply in the literature on microworld-based research on DDM. They also suggest specific changes to be made for the full-scale version of this study.

3.1 General considerations

Some considerations emerge from examining DDM from the point of view of system control theory. On that view, the objective of using microworlds to train DDM would be to impart some kind of non-proportional control strategy to the learner. Whether this would be an internal (mental) model-based continuous control strategy or a discrete, “bang-bang” strategy (perhaps based on appropriate decision heuristics) is as yet unclear, and further research is required to determine the best control strategy for real-time, long timescale DDM. However, in discussing the control strategies that participants seem to have used in this study, we noted that the natural tendency of the human controller to attempt proportional control in short lag systems is a major cause of performance problems in DDM. Thus, training interventions for DDM should not only focus on imparting the “correct” skills but also on learning to avoid the classic DDM errors due to the use of proportional control (such as inducing oscillations in systems dominated by negative feedback loops with delays). This could be achieved in a variety of ways. For instance, trainees could use microworlds to explore various behaviours other than the “correct” or desired ones in a microworld (essentially encouraging them to explore and make mistakes). Also, early stages of training could involve simple microworlds specially designed to expose typical DDM or systems thinking mistakes (e.g., the systems archetypes documented by researchers in the fields of system dynamics and management training, summarized in Senge [10], and applied to military DDM in Rehak et al. [8]).

Microworlds based on archetypes could also be used to apply a part-task training approach to DDM training. However, the usefulness of part-task training for DDM is difficult to assess a priori, since it implies two types of decomposition of DDM tasks, both of which require much further research and are beyond the scope of this report. Suffice it to say that the first type is a functional decomposition of the socio-technical systems that trainees are to learn to control into simpler systems. The second is a decomposition of the cognitive processes, mechanisms and sub-tasks that are required for performing DDM. The former topic is made complicated by the fact that the subcomponents of a complex dynamic system typically interact in non-linear fashion, and raises the issue of whether the dynamics of isolated subsystems can still be identified and recognized once the whole system has been reconstructed. The latter topic is made complicated

by the fact that, as noted by Rigas et al. [30], researchers have been unable to find strong correlations between typical psychometric measures (e.g., IQ tests) and DDM performance, and have to date only found weak correlations with the Raven's Progressive Matrices tests, which are supposed to index fluid intelligence and pattern-matching skills.

The present study suggests that the tempo of the training sessions might influence the control strategies used by trainees, and more generally that the tempo of the simulation might affect how people approach DDM tasks in general. On one hand, compressing time in a microworld might help trainees avoid discrete, all-or-nothing control and adopt continuous control strategies based on mental representations of the system. On the other hand, in a review of the role of time in decision making, Ariely and Zakay [31] note that time stress from task acceleration can change the way decisions are made, and even impair learning under some circumstances. Further research is therefore needed to examine how accelerating DDM tasks (and the degree of acceleration) affects the decision making task and whether these changes would be beneficial or detrimental to post-training DDM performance.

It is likely that presenting decision-support type information aimed at enhancing either the formation of a mental representation of the system or the use of appropriate heuristics would improve learning in microworlds. This could include: showing the structure of the model underlying the microworld, and the flow of resources and information through the model as the simulation unfolds (the so-called "glass box" approach trialed by Bondanella et. al. [21] for the US Army); displaying information in the microworld's GUI about internal variables of the model instead of only showing information about the input and output variables of the system (specifically, the variables that the decision maker directly manipulates and is interested in controlling respectively); a graph of the evolution of key system variable over time (known as a state diagram); and so on. However, it would be important when applying such training supports to distinguish between additional information that would support transfer and post-training performance rather than simply supporting performance in the training phase. Schmidt and Bjork [32] warn against designing training interventions that support performance during training without engaging skills that are required for generalizing performance beyond the training environment. For instance, displaying a visual warning about animosity levels in Region A might enhance performance without encouraging participants to elaborate their own mental representation of the dynamics of animosity and security in the microworld, and might hinder participants performing the task without the benefit of such an indicator in the transfer phase of the study or in the real world. On the other hand, displaying the evolution of key variables of the microworld over time might allow participants to elaborate mental models of the system; furthermore, the static record of the system's state over time at the end of a trial might encourage the development of speed-independent models of the system dynamics. Displaying such a graph on only a subset of training trials might further encourage participants to elaborate mental models in the absence of the supporting graph. More research is required to test these hypotheses and determine how to enhance both training performance and transfer in microworlds.

To conclude this discussion of general considerations on training with microworlds, it is important to note that military decision making, in particular at the strategic and operational levels, is inherently collective decision making (Brehmer [1], Clancy et al. [3]), often involving a variety of non-military governmental and non-governmental stakeholders in addition to the commander's personnel (Essens et al. [33]), and that any training intervention for military DDM will eventually have to address the group or team aspects of DDM. Collective DDM has not been extensively researched with microworlds, although Clancy et al. [3] have successfully used

microworlds to examine the effects of command style on team DDM. Future research is needed to address the team aspects of training DDM with microworlds. Specifically, what are the challenges of integrating individuals trained on faster-than-real-time microworlds into a functional team? Can team processes benefit from microworld training, and if so, which ones (communications, collective situation awareness, shared mental models, collective decision making, etc.)? Bakken and Gilliam [7] have suggested that microworlds are appropriate for learning the dynamics of the complex socio-technical system that military trainees will face in operational settings, but that the actual command staff procedures (reporting protocols, communications standards, etc.) need to be trained in real-time, high-fidelity settings. Assuming this hypothesis is borne out by future research, it would still be pertinent to study whether team members would benefit from training with microworlds that represent the dynamics of information flow through the team, and whether special procedures are required to ensure that individuals trained on microworlds are able to coalesce into a team.

3.2 Recommendations for improving the present study.

As discussed above, participants in the present pilot study performed sub-optimally despite showing some evidence of learning in the training phase. Also, it was found that the quantitative measures selected as the dependent variables in the study reflected some but not all of the noteworthy qualitative behavioural patterns displayed by the participants. Furthermore, the “real time” training condition was not an optimal control condition for the study as it consisted of a single presentation of each microworld model, therefore making it difficult to distinguish between possible benefits of accelerated microworld presentation on the perception of temporal dynamics and those of the extra rehearsals afforded by compressing time.

The participants in the pilot study expressed the opinion that the number of practice trials provided was insufficient for understanding the dynamics of the different models. Future versions of this experiment will include more training trials with each microworld models. Furthermore, participants should be encouraged to spend at least part of the training session exploring the dynamics of each microworld rather than just attempting to apply the “correct” decision strategy. There is much research that shows that learning negative instances (i.e., learning what doesn’t work well) can be just as important as learning positive instances (Gonzalez et al. [22] have applied this idea to studying learning in microworlds). Similarly, Miller et al. [20] suggest that learning in microworlds is improved when participants are allowed to explore the microworld without having certain learning goals prescribed to them. However, the same study indicates that providing learners with specific learning goals (e.g., learning to apply the right resource movement strategy) that support specific desired skills or curriculum items can be beneficial as well – so long as it is possible to identify the required skills. Given the lack of knowledge on the specific cognitive skills required for successful DDM, future versions of this study should use a combination of trials with specific task instructions of the type used in the pilot study, and trials encouraging exploration of the microworld models.

The provision of additional information (e.g., displaying more internal model variables) to support DDM is a complex issue in and of itself and is best addressed in a dedicated study. Nevertheless, participants in the study felt the cover story and suggestions displayed on the screen should be more salient. This could be achieved by separating the text and the microworld GUI itself, and by designing the display so that the participants must read through the text before starting the simulation.

Further research is required for selecting numerical measures that better capture the qualitative patterns observed in individual participants' runs. As noted above, the average movement size (average number of units moved per movement) is a simple measure that could prove to be useful. Another similar measure is the number of changes in movement direction (i.e., the number of sign changes observed in the resource movement variable). Analysis of the present study was restricted to measures that were easily analyzed using STATISTICA (Statsoft Inc.), that is, aggregate measures that could be easily obtained and submitted to statistical significance tests. Other quantitative methods that are better performed with other software but that might better capture temporal patterns (e.g., Fourier analysis) or that might characterize the "dimension reduction" (variable aggregation) that participants might have performed during the task (e.g., by applying Latent Semantic Analysis to microworld performance logs as proposed by Quesada, Kintsch & Gomez [34]) should be considered in the future.

The individual differences in performance that were observed in this study (i.e., those with previous simulation game experience performed better) strongly suggest that participants should be carefully selected and assigned to conditions in DDM research based on their pre-training DDM abilities. In addition to self-report or questionnaire-based assessment of simulation game experience, it might be possible to sort participants based on more objective measures. Rigas et al. [30] found moderate correlations between DDM performance and Raven's Progressive Matrices (a test of "fluid" intelligence and pattern recognition abilities), and Lerch & Harter [16] report positive effects of individual working memory capacity on DDM performance in the DDM literature (unexpectedly their own study showed performance decrements associated with high working memory capacity). Participants should be assessed with a combination of such techniques prior to the DDM study. They could then be assigned to different ability groups to test the effects of these individual differences on training DDM, or simply assigned in a balanced manner to both training conditions in order to mitigate the effects of individual differences on the study.

Finally, in order to disentangle the effects of training simulation speed and number of extra repetitions afforded by a time-compressed microworld, a different control condition is needed from the one used in the pilot study. Specifically, it would be beneficial to compare the performance of a group trained at the same speed as the accelerated group in the pilot study (where one simulated minute lasted 3 real seconds) with a group performing the same number of training trials at a slower speed (e.g., where one minute of simulated time would last 30 seconds of real time), while keeping the transfer session at the same speed as in this study. By keeping the number of training trials constant but varying the training speed between groups, it will be possible to determine whether any positive training effects of accelerated training are due to time compression making the dynamics of the microworld more salient to the trainees, or whether they are simple due to the extra training trials that can take place when simulation time is compressed. Such a design would also allow the examination of the alternative hypothesis that recognition of dynamic patterns degrades with larger differences between the speed at which the patterns were acquired and the speed at which the patterns occur in real time. Future research could also examine the combination of various training speeds to facilitate tempo-independent recognition and prevent the fixation on specific control strategies that particular simulation speeds seem to have induced in this pilot study.

4. Conclusion

The present pilot study aimed at evaluating an experimental procedure for examining the transfer of DDM skills, specifically, the dynamic pattern recognition that is thought to be necessary for DDM, from an accelerated microworld training environment to a “real-time” target environment. The pilot study showed that the training procedure that was designed for the experiment is insufficient and needs to be improved. The study also showed, however, some evidence of learning in the training sessions and some evidence of transfer of this learning to the target environment. The study helped to determine that the numeric measures selected for the study capture some but not all of the relevant qualitative patterns that were observed in the data from the participants, and suggested additional measures to be considered in the future. It was also found that the control condition used in the study did not allow a distinction between the possible benefits of accelerated training due to enhanced dynamic pattern salience and those due to extra rehearsal.

Based on the findings of the pilot study and on the research literatures on training and decision-making, a number of considerations for improving training with microworlds in general were identified, and a number of recommendations to improve the experimental design used in this pilot study were proposed. The main considerations are summarized here:

- The cognitive skills and control processes that underlie particular DDM tasks must be clearly identified and assessed to ensure that they are supported by the design of the microworld used for training the task.
- Training for DDM should target both specific skills for the target task and generic “systems thinking” skills to avoid common DDM “mistakes”; to this end microworlds should be used to learn system dynamics through exploration and trial-and-error, and not just to learn the “correct” procedures.
- Microworld design should support not only the development of skills during training but also the transfer of these skills to the real-time, post-training environment.
- In particular, microworlds should support the application of DDM skills independently of the real-time speed of the dynamics being learned; this could include speed-independent representations of the dynamics of the system, or presenting microworlds at a variety of different speeds.
- The real-time speed of training likely influences the control strategies learned in training, and its effects on post-training performance must therefore be further studied.
- The effects of microworld training on team coordination and performance need to be examined in more depth, in particular with respect to the effect of time compression on team dynamics, and to the use of microworlds to learn the dynamics of a particular team.

References

- [1] Brehmer, B. (1992). Dynamic decision making: human control of complex systems. *Acta Psychologica*, 81, 211-241.
- [2] Brehmer, B. (1995). Feedback delays in complex dynamic decision tasks. In P.A. Frensch and J. Funke (Eds.), *Complex problem solving: the European perspective*, pp. 103-130. Mahwah, NJ: Lawrence Erlbaum Associates.
- [3] Clancy, J. M., Elliott, G. C., Ley, T., Omodei, M. M., Wearing, A. J., McLennan, J. & Thorsteinsson, E. B. (2003). Command style and team performance in dynamic decision-making tasks. In S. L. Schneider & J. Shanteau (Eds.), *Emerging perspectives on judgment and decision research*, pp. 586-619. New York: Oxford University Press.
- [4] Dörner, D. (1996). The logic of failure: recognizing and avoiding error in complex situations. New York: Metropolitan Books.
- [5] Sterman, J.D. (1994). Learning in and about complex systems. *System Dynamics Review*, 10, 291-330.
- [6] Bakken, B. T., & Vamraak, T. (2003). Misperception of dynamics in military planning: Exploring the counter-intuitive behaviour of the logistics chain. In *Proceedings of the 21st International Conference of the System Dynamics Society*. Albany, NY: System Dynamics Society.
- [7] Bakken, B.T., & Gilljam, M. (2003). Training to improve decision making – system dynamics applied to higher-level military operations. *Journal of Battlefield Technology*, 6, 33-42.
- [8] Rehak, L. A., Lamoureux, T. M., & Bos, J. C. (2006). *Systems archetypes for military dynamic decision making*. (DRDC Toronto CR2006-202). Defence Research & Development Canada—Toronto.
- [9] Haberstroh, S., Betsch, T., Glöckner, A., Haar, T., & Stiller, A. (2005). The impact of routines on deliberate decisions: the microworld-simulation COMMERCE. In T. Betsch & S. Haberstroh (Eds.), *The routines of decision making*, pp. 211-230. Mahwah, NJ: Lawrence Erlbaum Associates.
- [10] Senge, P. M. (1990). The fifth discipline: the art & practice of the learning organization. New York: Doubleday.
- [11] Shanteau, J., Friel, B. M., Thomas, R. R., & Raacke, J. (2005). Development of expertise in a dynamic decision-making environment. In T. Betsch & S. Haberstroh (Eds.), *The routines of decision making*, pp. 251-270. Mahwah, NJ: Lawrence Erlbaum Associates.
- [12] Gonzalez, C., Vanyukov, P., & Martin, M. K. (2005). The use of microworlds to study dynamic decision making. *Computers in Human Behavior*, 21, 273-286.

- [13] Rouwette, E. A. J. A., Größler, A., & Vennix, J. A. M. (2004). Exploring influencing factors on rationality: A literature review of dynamic decision-making studies in system dynamics. *Systems Research and Behavioral Science*, 21, 351-370.
- [14] Gonzalez, C. (2005). Decision support for real-time, dynamic decision-making tasks. *Organizational Behavior and Human Decision Processes*, 96, 142-154.
- [15] Gonzalez, C., & Quesada, J. (2003). Learning in dynamic decision making: The recognition process. *Computational & Mathematical Organization Theory*, 9, 287-304.
- [16] Lerch, F. J., & Harter, D. E. (2001). Cognitive support for real-time dynamic decision making. *Information Systems Research*, 12, 63-82.
- [17] Ali, S.F., Guckenberger, D., Rossi, M., & Williams, M. (2000, May). Evaluation of above real-time training and self-instructional strategies for airmanship tasks on a flight simulator (AFRL-HE-AZ-TR-2000-0112, ADA391561). Proj 1123. F41624-98-1-0005, Tuskegee University, AL. DTIC.
- [18] Levy, D. G., Lewis, M. W., Bondanella, J. R., Baisden, M., & Ettinger, E. (2001). *Exploring the use of microworld models to train Army logistics management skills*. Santa Monica, CA: RAND.
- [19] Martin, L. M. W., Shirley, M., & McGinnis, M. (1988). Microworlds to macroworlds: An experiment in the conceptual transfer of ecological concepts. *Children's Environments Quarterly*, 5(4), 32-38
- [20] Miller, C. S., Lehman, J. F., & Koedinger, K. R. (1999). Goals and learning in microworlds. *Cognitive Science*, 23, 305-336.
- [21] Bondanella, J. R., Lewis, M. W., Steinberg, P. S., Park, G. S., Levy, D. G., Ettinger, E., Oaks, D. M., Sollinger, J. M., Winkler, J. D., Halliday, J. M., & Way-Smith, S. (1998). *Microworld simulations for command and control training of theater logistics and support staffs: A curriculum strategy*. Santa Monica, CA: RAND.
- [22] Gonzalez, C., Lerch, J. F., Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, 27, 591-635.
- [23] Klein, G. (1989). Recognition-primed decisions. In W. B. Rouse (Ed.), *Advances in man-machine systems research*, Vol. 5, pp. 47-92. Greenwich, CT: JAI Press.
- [24] Klein, G. (1997). The recognition-primed decision (RPD) model: looking back, looking forward. In C. E. Zsombok & G. Klein (Eds.), *Naturalistic decision making*, pp. 285-292. Mahwah, NJ: Lawrence Erlbaum Associates.
- [25] Boltz, M.G. (1998). The processing of temporal and nontemporal information in the remembering of event durations and musical structure. *J. Exp. Psy.: Hum. Perc. & Perf.*, 24, 1087-1104.
- [26] Schulkind, M.D. (1999). Long-term memory for temporal structure: Evidence from the identification of well-known and novel songs. *Memory & Cognition*, 27, 896-906.

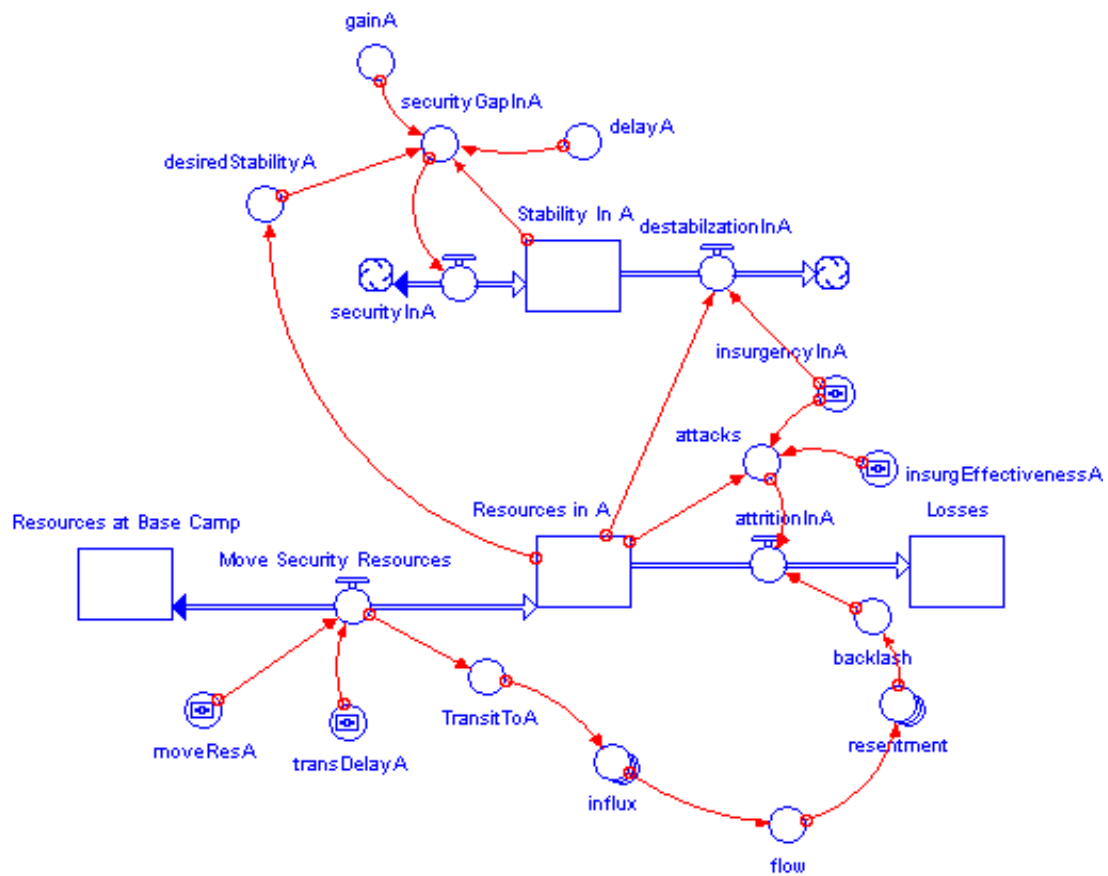
- [27] DRDC Toronto (2003). *DRDC Guidelines for Human Subject Participation in Research Projects*. Defence Research and Development Canada – Toronto.
- [28] Pigeau, R. (2005). Command Effectiveness and Behaviour Section guide to stress compensation for human subjects. Defence Research and Development Canada – Toronto.
- [29] Jagacinski, R. J. & Flach, J. M. (2003). Control theory for humans: quantitative approaches to modeling performance. Mahwah, NJ: Erlbaum.
- [30] Rigas, G., Carling, E., & Brehmer, B. (2002). Reliability and validity of performance measures in microworlds. *Intelligence*, 30, 463-480.
- [31] Ariely, D., & Zakay, D. (2001). A timely account of the role of duration in decision making. *Acta Psychologica*, 108, 187-207.
- [32] Schmidt, R.A., & Bjork, R.A. (1992). New conceptualizations of practice: common principles in three paradigms suggest new concepts for training. *Psychological Science*, 3, 207-217.
- [33] Essens, P., Vogelaar, A., Mylle, J., Blendell, C., Paris, C., Halpin, S., & Baranski, J. (2005). Military command team effectiveness: model and instrument for assessment and improvement. (AC/323(HFM-087)TP/59). NATO Research and Technology Organization.
- [34] Quesada, J., Kintsch, W., & Gomez, E. (in press). A computational theory of complex problem solving using Latent Semantic Analysis. Submitted to *Cognitive Science*.

This page intentionally left blank.

Annex A Microworld Models and Equation Listings

A.1 Model A1

A.1.1 Model Diagram

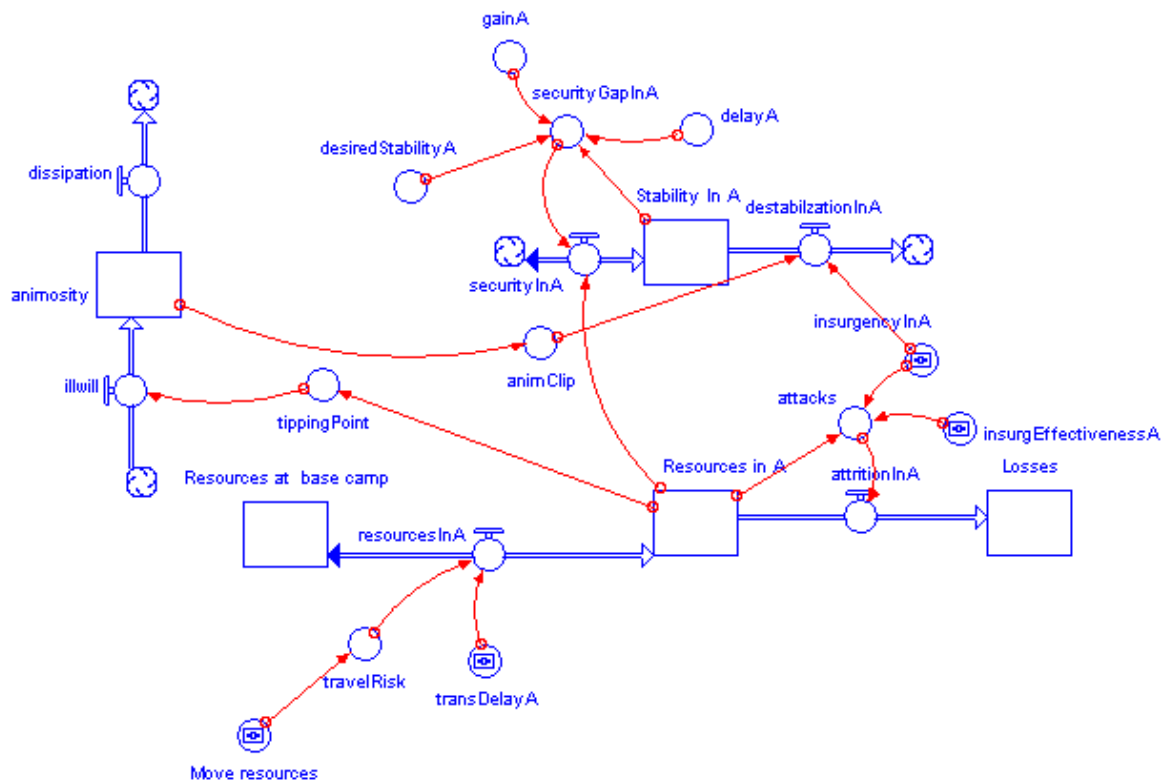


A.1.2 Equation listing

```
Losses(t) = Losses(t - dt) + (attritionInA) * dt
INIT Losses = 0
attritionInA = attacks+backlash
Resources_at_Base_Camp(t) = Resources_at_Base_Camp(t - dt) + (-
Move_Security_Resources) * dt
INIT Resources_at_Base_Camp = 75
Move_Security_Resources = DELAY(moveResA,transDelayA, moveResA)
Resources_in_A(t) = Resources_in_A(t - dt) + (Move_Security_Resources - attritionInA) * dt
INIT Resources_in_A = 0
Move_Security_Resources = DELAY(moveResA,transDelayA, moveResA)
attritionInA = attacks+backlash
Stability_In_A(t) = Stability_In_A(t - dt) + (securityInA - destabilizationInA) * dt
INIT Stability_In_A = 5
securityInA = securityGapInA + RANDOM(0,1)
destabilizationInA = IF (Resources_in_A < 20) OR (Resources_in_A > 23) THEN (insurgencyInA
+ Random(0,1)) ELSE RANDOM(0,1)
attacks = IF (Resources_in_A > 19) THEN (insurgencyInA*insurgEffectivenessA * 2) ELSE
(insurgencyInA*insurgEffectivenessA)
backlash = IF(ARRAYSUM(resentment[*]) > 10) THEN (.25 * ARRAYMEAN(resentment[*]))
ELSE 0
clock = INT(TIME / 4) + (MOD(INT(TIME / DT), 60) / 100)
Days = TIME
delayA = 1
desiredStabilityA = Resources_in_A
flow = ARRAYSUM(influx[*])
gainA = .4
influx[1] = TransitToA
influx[2] = DELAY(TransitToA,1,TransitToA)
influx[3] = DELAY(TransitToA,2,TransitToA)
insurgEffectivenessA = .2
insurgencyInA = .3
moveResA = 0
resentment[1] = DELAY(flow, 2, flow)
resentment[2] = DELAY(flow, 4, flow)
resentment[3] = DELAY(flow, 6, flow)
securityGapInA = gainA * (desiredStabilityA - DELAY(Stability_In_A,delayA,Stability_In_A))
transDelayA = 0
TransitToA = Move_Security_Resources
```

A.2 Model A2

A.2.1 Model diagram

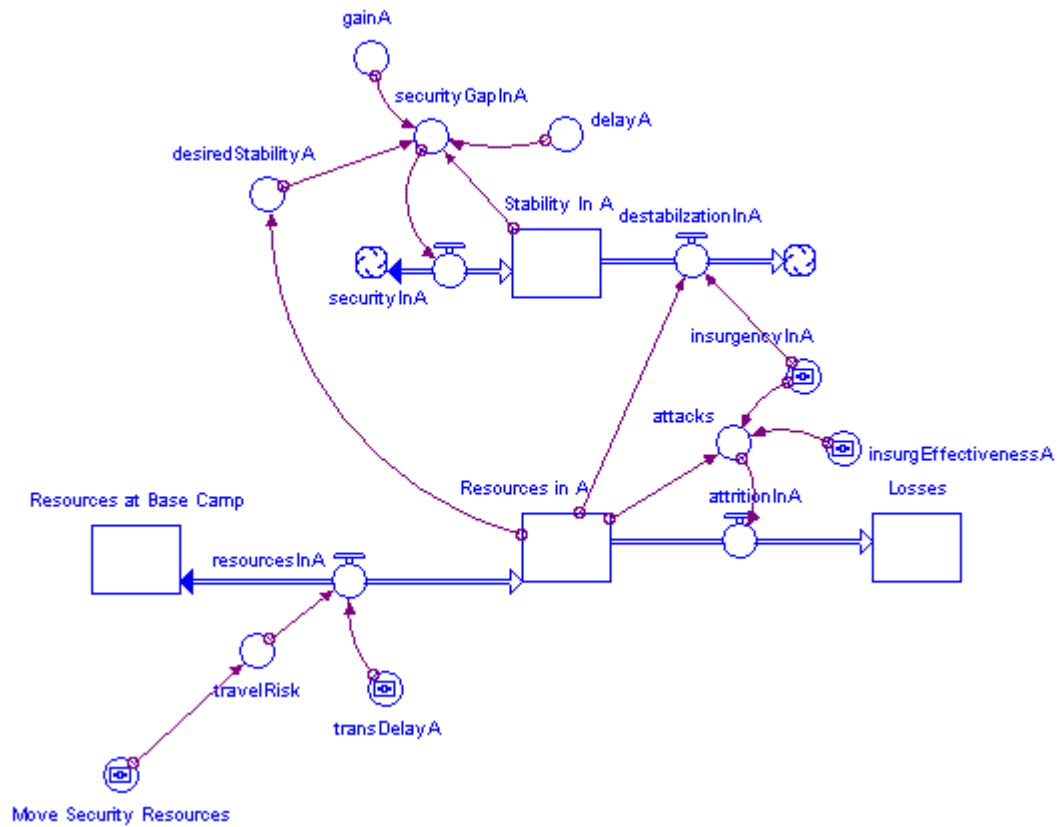


A.2.2 Equation listing

```
Losses(t) = Losses(t - dt) + (attritionInA) * dt
INIT Losses = 0
attritionInA = attacks+backlash
Resources_at_Base_Camp(t) = Resources_at_Base_Camp(t - dt) + (-
Move_Security_Resources) * dt
INIT Resources_at_Base_Camp = 75
Move_Security_Resources = DELAY(moveResA,transDelayA, moveResA)
Resources_in_A(t) = Resources_in_A(t - dt) + (Move_Security_Resources - attritionInA) * dt
INIT Resources_in_A = 0
Move_Security_Resources = DELAY(moveResA,transDelayA, moveResA)
attritionInA = attacks+backlash
Stability_In_A(t) = Stability_In_A(t - dt) + (securityInA - destabilizationInA) * dt
INIT Stability_In_A = 5
securityInA = securityGapInA + RANDOM(0,1)
destabilizationInA = IF (Resources_in_A < 20) OR (Resources_in_A > 23) THEN (insurgencyInA
+ Random(0,1)) ELSE RANDOM(0,1)
attacks = IF (Resources_in_A > 19) THEN (insurgencyInA*insurgEffectivenessA * 2) ELSE
(insurgencyInA*insurgEffectivenessA)
backlash = IF(ARRAYSUM(resentment[*]) > 10) THEN (.25 * ARRAYMEAN(resentment[*]))
ELSE 0
clock = INT(TIME / 4) + (MOD(INT(TIME / DT), 60) / 100)
Days = TIME
delayA = 1
desiredStabilityA = Resources_in_A
flow = ARRAYSUM(influx[*])
gainA = .4
influx[1] = TransitToA
influx[2] = DELAY(TransitToA,1,TransitToA)
influx[3] = DELAY(TransitToA,2,TransitToA)
insurgEffectivenessA = .2
insurgencyInA = .3
moveResA = 0
resentment[1] = DELAY(flow, 2, flow)
resentment[2] = DELAY(flow, 4, flow)
resentment[3] = DELAY(flow, 6, flow)
securityGapInA = gainA * (desiredStabilityA - DELAY(Stability_In_A,delayA,Stability_In_A))
transDelayA = 0
TransitToA = Move_Security_Resources
```

A.3 Model B1

A.3.1 Model diagram



A.3.2 Equation listing

```
Losses(t) = Losses(t - dt) + (attritionInA) * dt
INIT Losses = 0
attritionInA = attacks
Resources_at_Base_Camp(t) = Resources_at_Base_Camp(t - dt) + (- resourcesInA) * dt
INIT Resources_at_Base_Camp = 75
resourcesInA = DELAY(travelRisk, transDelayA, travelRisk)
Resources_in_A(t) = Resources_in_A(t - dt) + (resourcesInA - attritionInA) * dt
INIT Resources_in_A = 0
resourcesInA = DELAY(travelRisk, transDelayA, travelRisk)
attritionInA = attacks
Stability_In_A(t) = Stability_In_A(t - dt) + (securityInA - destabilizationInA) * dt
INIT Stability_In_A = 5
securityInA = securityGapInA + RANDOM(0,1)
destabilizationInA = IF (Resources_in_A < 20) OR (Resources_in_A > 23) THEN (insurgencyInA
+ Random(0,1)) ELSE RANDOM(0,1)
attacks = IF (Resources_in_A < 17) THEN (insurgencyInA*insurgEffectivenessA * 6) ELSE
(insurgencyInA*insurgEffectivenessA)
clock = INT(TIME / 4) + (MOD(INT(TIME / DT), 60) / 100)
Days = TIME
delayA = 1
desiredStabilityA = Resources_in_A
gainA = .4
insurgEffectivenessA = .5
insurgencyInA = .3
Move_Security_Resources = 0
securityGapInA = gainA * (desiredStabilityA - DELAY(Stability_In_A,delayA,Stability_In_A))
transDelayA = 1
travelRisk = IF (Move_Security_Resources >= -3) AND (Move_Security_Resources <= 3) THEN
(RANDOM(0, 1) * Move_Security_Resources) ELSE Move_Security_Resources
```


A.4.2 Equation listing

```
animosity(t) = animosity(t - dt) + (illwill - dissipation) * dt
  INIT animosity = 0
illwill = tippingPoint
dissipation = .7
Losses(t) = Losses(t - dt) + (attritionInA) * dt
  INIT Losses = 0
attritionInA = attacks
Resources_at__base_camp(t) = Resources_at__base_camp(t - dt) + (- resourcesInA) * dt
  INIT Resources_at__base_camp = 75
resourcesInA = DELAY(travelRisk, transDelayA, travelRisk)
Resources_in_A(t) = Resources_in_A(t - dt) + (resourcesInA - attritionInA) * dt
  INIT Resources_in_A = 0
resourcesInA = DELAY(travelRisk, transDelayA, travelRisk)
attritionInA = attacks
Stability_In_A(t) = Stability_In_A(t - dt) + (securityInA - destabilizationInA) * dt
  INIT Stability_In_A = 5
securityInA = IF (Resources_in_A > 15) THEN (securityGapInA + RANDOM(0,1)) ELSE
RANDOM(0,1)
destabilizationInA = insurgencyInA * animClip + Random(0,1)
animClip = IF (animosity < 4) THEN .6 ELSE animosity
attacks = IF (Resources_in_A > 19) THEN (insurgencyInA*insurgEffectivenessA * 2) ELSE
(insurgencyInA*insurgEffectivenessA)
clock = INT(TIME / 4) + (MOD(INT(TIME / DT), 60) / 100)
Day = TIME
delayA = 0
desiredStabilityA = 18
gainA = .4
insurgEffectivenessA = .2
insurgencyInA = .3
Move_resources = 0
securityGapInA = gainA * (desiredStabilityA - DELAY(Stability_In_A,delayA,Stability_In_A))
tippingPoint = Resources_in_A / 17
transDelayA = 3
travelRisk = IF (Move_resources >= -3) AND (Move_resources <= 3) THEN (RANDOM(0, 1) *
Move_resources) ELSE Move_resources
```

This page intentionally left blank.

List of symbols/abbreviations/acronyms/initialisms

ANOVA	Analysis of Variance
CF	Canadian Forces
DDM	Dynamic Decision Making
DND	Department of National Defence
DRDC	Defence Research and Development Canada
GUI	Graphical User Interface
OPI	Office of Primary Interest
R&D	Research & Development
RT	Real-time

This page intentionally left blank.

Distribution list

Document No.: DRDC Toronto TR 2006-239

LIST PART 1: Internal Distribution by Centre:

2 – Library

TOTAL LIST PART 1: 2

LIST PART 2: External Distribution by DRDKIM

1 – Defence R&D Canada – Centre for Operational Research and Analysis

4 – Defence R&D Canada – Valcartier

Attn: P. Lessard

R. Breton

M.E. Jobidon

M. Bélanger

2 – Director General Land Capabilities Development

Attn: R. Reshke

LCol R.W. Bell

2 – Canadian Forces Experimentation Centre

Attn: F. Lichacz

P. Farrell

1 – Directorate of Army Training

Attn: Maj C. Horeczy

TOTAL LIST PART 2: 10

TOTAL COPIES REQUIRED: 12

This page intentionally left blank.

UNCLASSIFIED

DOCUMENT CONTROL DATA (Security classification of the title, body of abstract and indexing annotation must be entered when the overall document is classified)		
1. ORIGINATOR (The name and address of the organization preparing the document, Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's document, or tasking agency, are entered in section 8.) Publishing: DRDC Toronto Performing: DRDC Toronto Monitoring: Contracting:		2. SECURITY CLASSIFICATION (Overall security classification of the document including special warning terms if applicable.) UNCLASSIFIED
3. TITLE (The complete document title as indicated on the title page. Its classification is indicated by the appropriate abbreviation (S, C, R, or U) in parenthesis at the end of the title) Accelerated Training for Command Dynamic Decision Making: A Pilot Study Using Microworlds (U) Formation accélérée de prise de décision dynamique pour le commandement : Étude pilote utilisant des micro-mondes		
4. AUTHORS (First name, middle initial and last name. If military, show rank, e.g. Maj. John E. Doe.) Jerzy Jarmasz		
5. DATE OF PUBLICATION (Month and year of publication of document.) December 2006	6a NO. OF PAGES (Total containing information, including Annexes, Appendices, etc.) 66	6b. NO. OF REFS (Total cited in document.) 34
7. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of document, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.) Technical Report		
8. SPONSORING ACTIVITY (The names of the department project office or laboratory sponsoring the research and development – include address.) Sponsoring: Tasking:		
9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant under which the document was written. Please specify whether project or grant.) Thrust 6b		9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)
10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document) DRDC Toronto TR2006–239		10b. OTHER DOCUMENT NO(s). (Any other numbers under which may be assigned this document either by the originator or by the sponsor.)
11. DOCUMENT AVAILABILITY (Any limitations on the dissemination of the document, other than those imposed by security classification.) Unlimited distribution		
12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, when further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected.) Unlimited announcement		

UNCLASSIFIED

UNCLASSIFIED

DOCUMENT CONTROL DATA

(Security classification of the title, body of abstract and indexing annotation must be entered when the overall document is classified)

13. **ABSTRACT** (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

(U) Dynamic Decision Making (DDM) is a skill that is both increasingly required and difficult to train for military commanders in today's security environment. Because it requires the timely sequencing of interdependent decisions in order to control complex and non-linear systems, DDM is a difficult skill to acquire for humans. Microworlds, stripped down simulations that focus on the dynamics of the target systems, have been proposed by many as training environments for DDM that avoid the time commitment, cost and personal danger of training command decision making with full-scale exercises or through mission experience. However, little research has been conducted on the factors that lead to effective microworld-based training. Specifically, it is unknown whether the time compression that occurs in microworlds enhances or inhibits the learning and transfer of complex system dynamics. A pilot study was conducted to examine whether participants are able to learn a simple DDM task in an accelerated microworld environment and then perform the same task in a similar but much slower environment. The results suggest that compressed-time microworlds can support training and transfer of DDM skills to "real-time" environments but that much remains to be learned about the conditions that favour the learning of DDM skills. Based on these results, general considerations for training DDM with microworlds and specific recommendations for improving the current study are provided.

14. **KEYWORDS, DESCRIPTORS or IDENTIFIERS** (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

(U) accelerated training; dynamic decision making; microworlds; command and control; command decision making

UNCLASSIFIED

Defence R&D Canada

Canada's Leader in Defence
and National Security
Science and Technology

R & D pour la défense Canada

Chef de file au Canada en matière
de science et de technologie pour
la défense et la sécurité nationale



www.drdc-rddc.gc.ca

